Christof Teuscher

Novel Computational Architectures for Nano-scale Devices

Los Alamos National Laboratory
Advanced Computing Laboratory (CCS-1)

www.teuscher.ch/christof

christof@lanl.gov


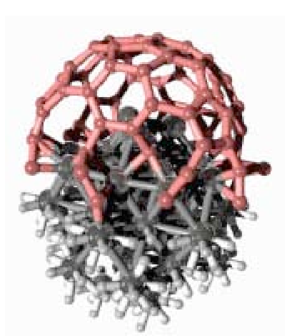
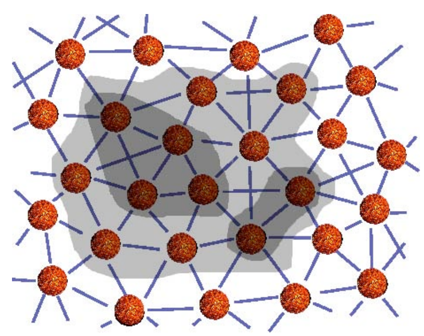
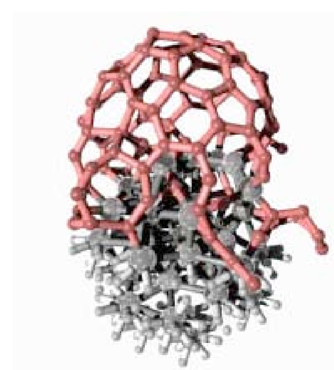
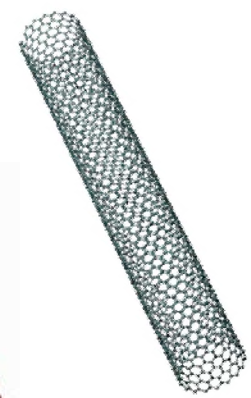

LA-UR-05-??

General Goals of this Talk

- Outline the challenges of building a nano-scale computational architecture.
- Present some of the major computational architectures and highlight weaknesses and strengths.
- Discuss our visions and directions of research.

- Provide global picture and avoid details.

Christof Teuscher • christof@lanl.gov

More More More More More More More More

Moore's Law

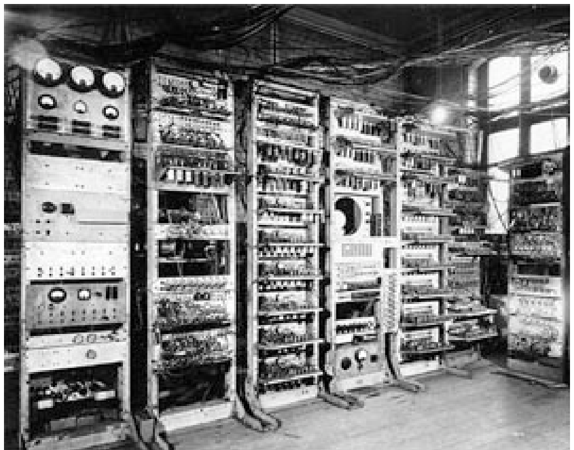

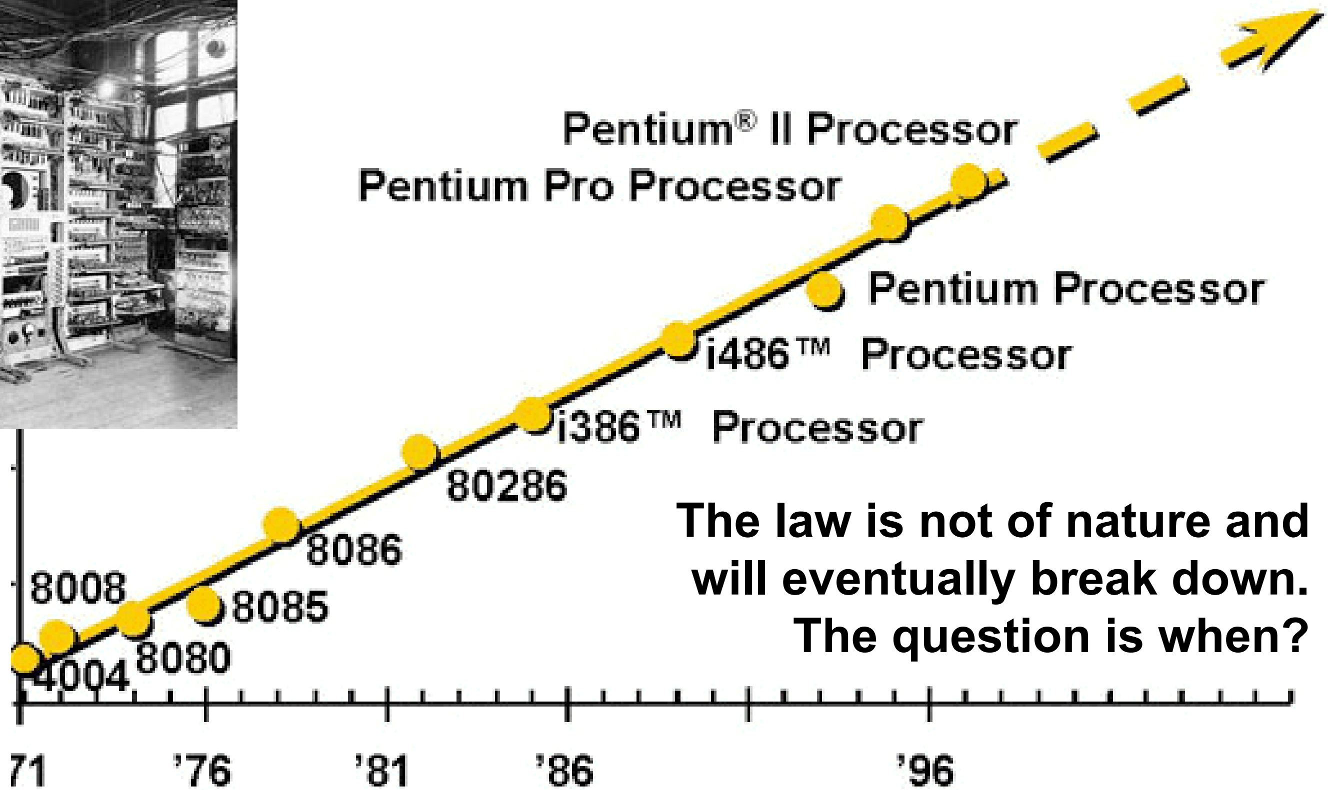

The law is not of nature and will eventually break down. The question is when?

Christof Teuscher • christof@lanl.gov

Smaller smaller smaller smaller smaller smaller smaller smaller smaller

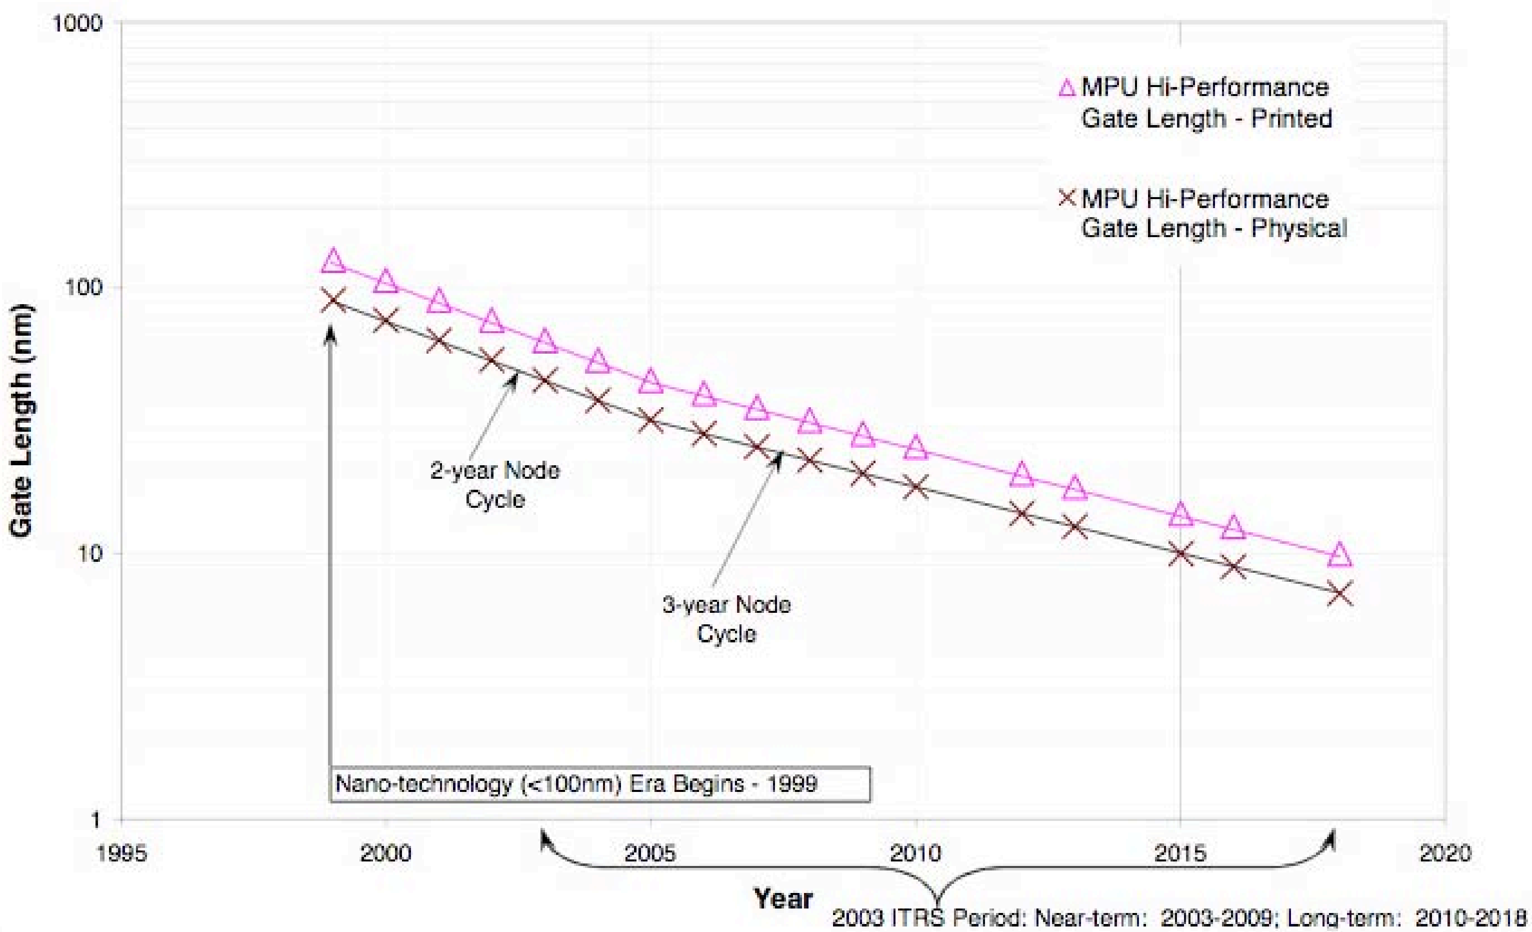

Christof Teuscher • christof@lanl.gov

Higher Higher Higher Higher Higher Higher Higher

Exponential Increase of Wafer Fabrication Costs

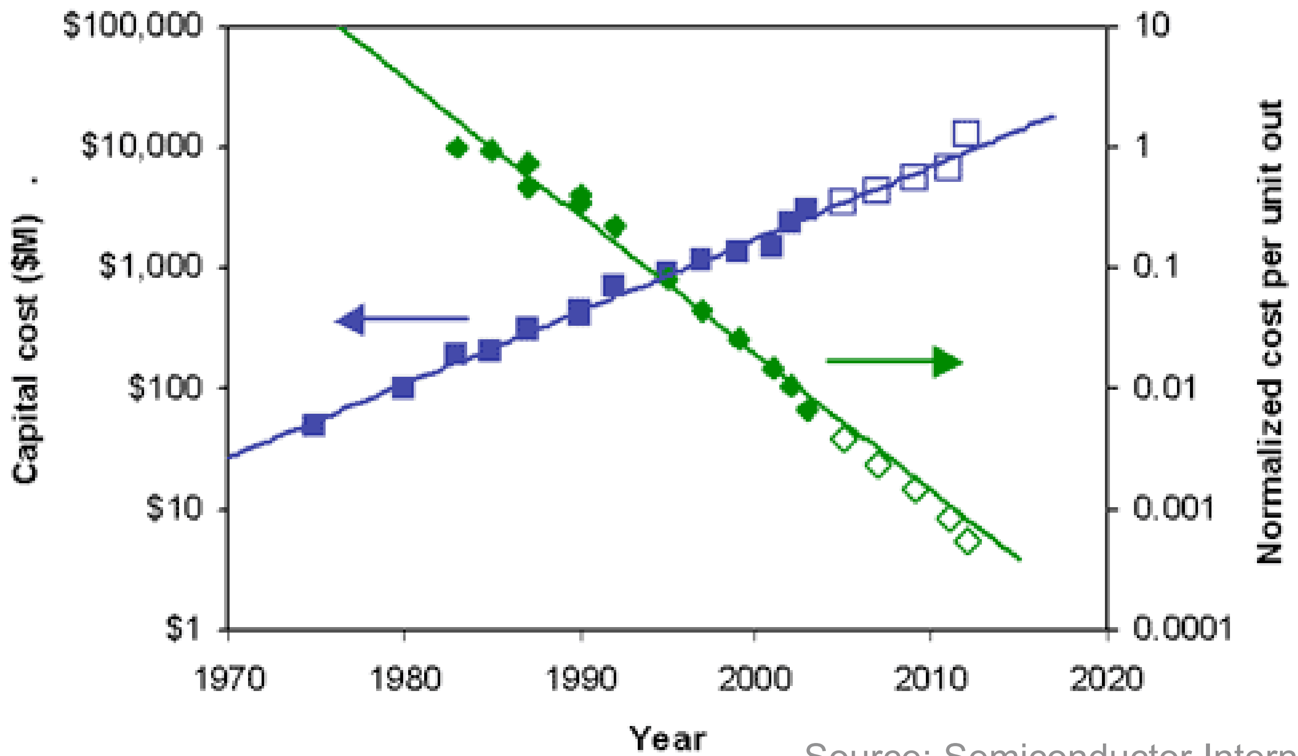

Source: Semiconductor International

The Productivity Design Gap

"Moore doesn't mean better!"

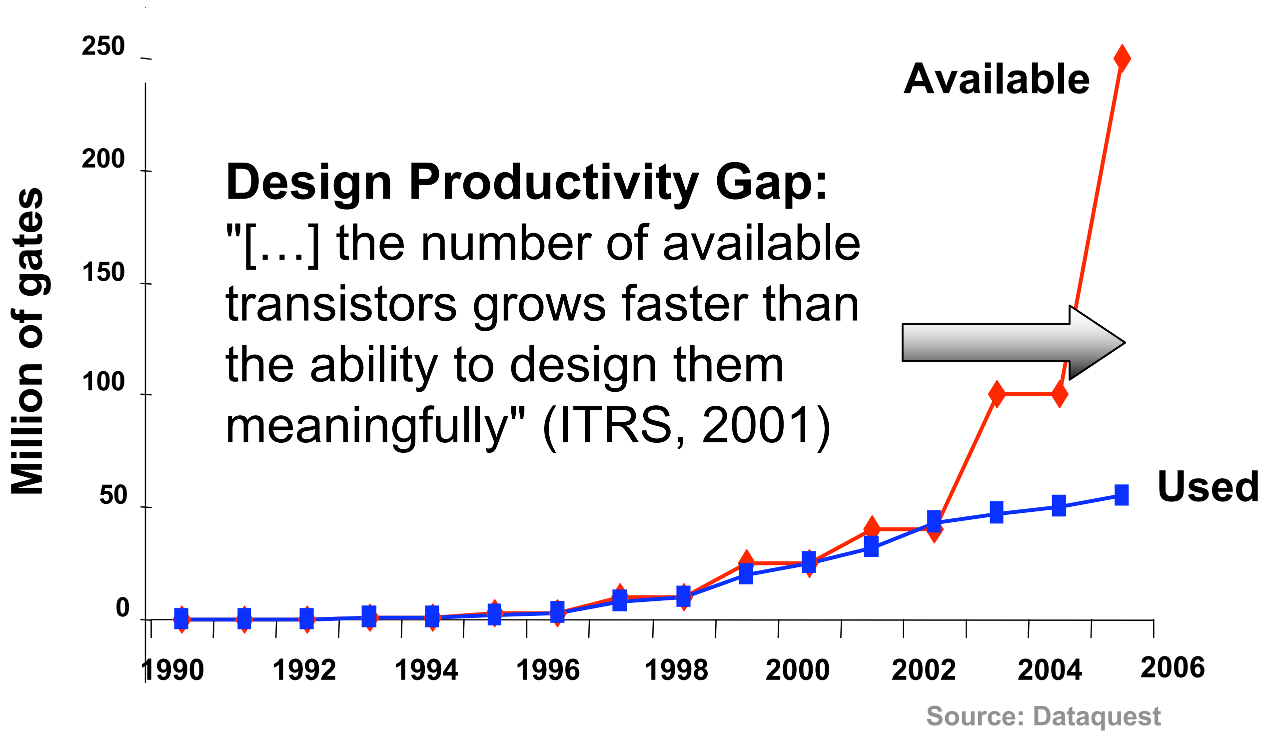

Christof Teuscher • christof@lanl.gov

What are the Long-term Goals of Computational Science?

- Compute faster, and faster, and faster and…
- Reduce energy consumption, area used, and cost.
- Maybe…make machines more "intelligent"

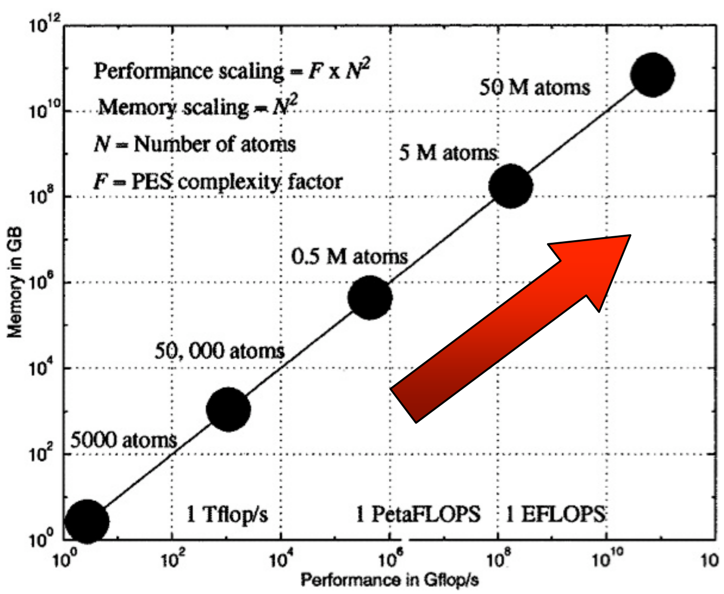

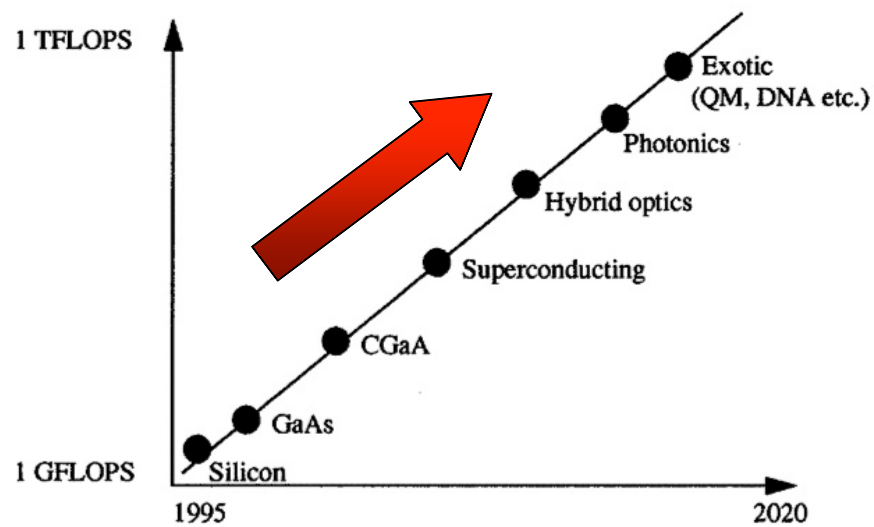

How can Nano-scale Devices Help us to Achieve this Goal?

- Miniaturization potentially helps to:
  - speed up computing
  - reduce energy consumption
  - reduce the area used
  - reduce fabrication costs, at least if self-assembly is used
  - increase parallelism
- But:
  - increases defect and fault rates
  - makes programming and configuration harder
- Ralph C. Merkle:
  - The objective is to design, model, and fabricate molecular machines (1991)

Short-term Goals

- Come up with a **computational architecture** as a proof-of-concept that has the potential to outperform traditional silicon approaches.
- **Scalability** is more important than speed, etc. in a first step.
- **But**: the opinions differ widely as to the type of architecture most suitable for achieving tremendous performance gains with computers built by nanotechnology.

Christof Teuscher • christof@lanl.gov

But…is Smaller Always Better?

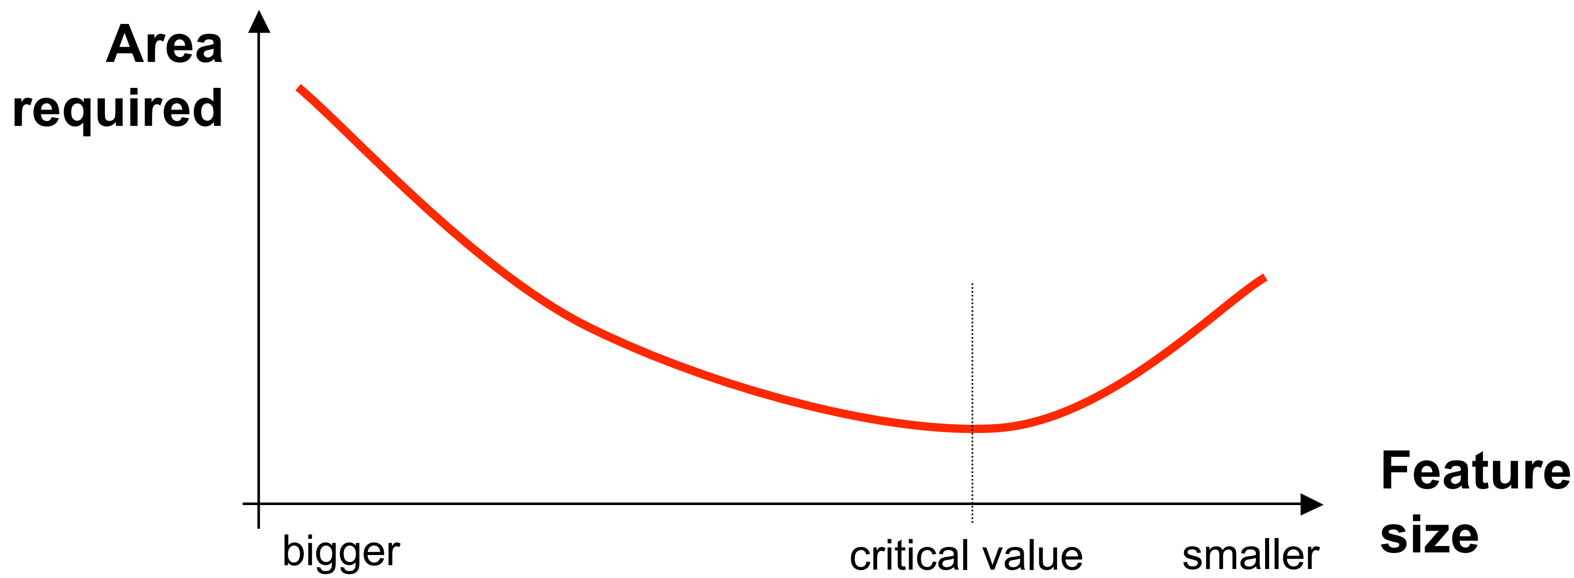

smaller → more unreliable (noise, etc.) → more redundancy needed → more area needed

Christof Teuscher • christof@lanl.gov

Three Things we Need to Make the Nano-Computing Revolution Happen

1. Invent a nanoscale device that **switches** electric current on and off. → nano transistor **done ☑**
2. Build a nanoscale circuit that controllably **links** very large numbers these devices together in order to perform memory and logic functions. **to do ☐**
3. Design an **architecture** (e.g., a micro processor) that allows the circuits to communicate with other systems independently of their lower-level details. **to do ☐**

Chen et al., Nanotechnology 14 (2003) 462-468.

Christof Teuscher • christof@lanl.gov

From the Device to the Architecture

The realization of a single working device is not a proof that many of them can be combined to form a computational circuit!

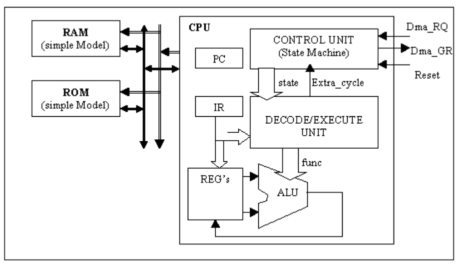

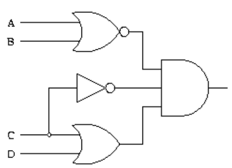

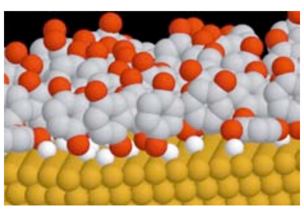

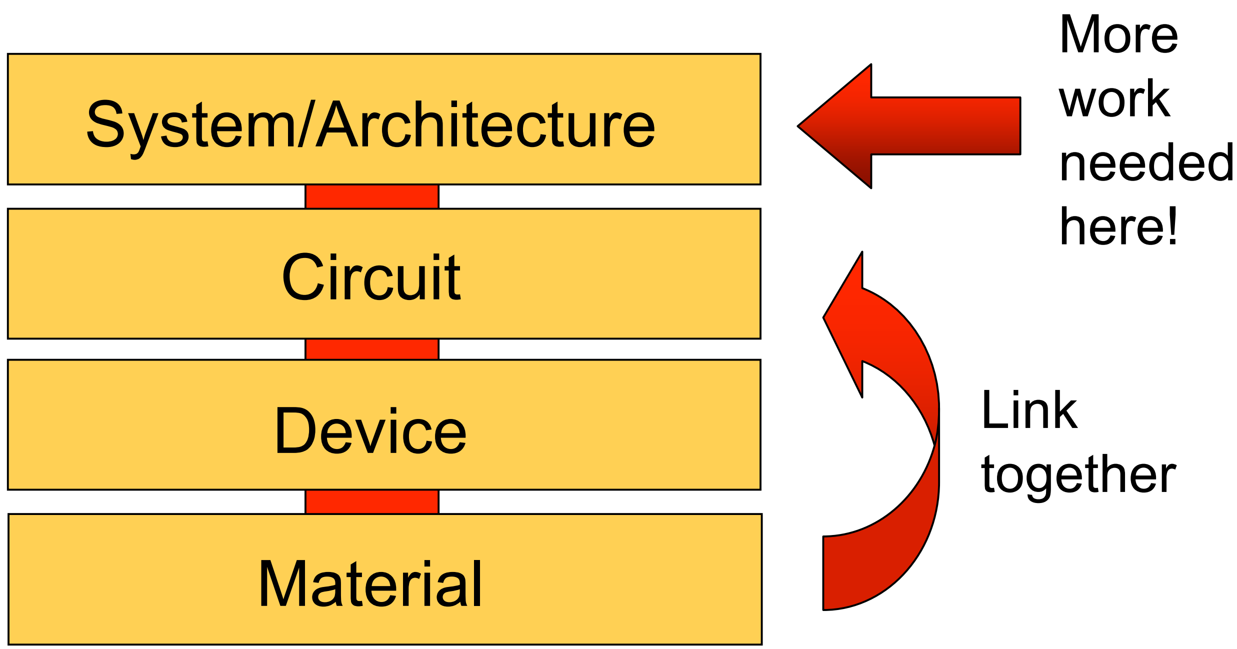

Some of the Big Challenges…

- We need to deal with:
  - **Unreliability** → up to 15% defects, unknown fault rate
  - **Billions** of components
  - **Inhomogeneous** substrate due to fabrication, e.g. self-assembly
  - **Asynchrony**, as global signals and synchronization become impossible
- Solutions to these challenges involve:
  - **Defect and fault tolerance** (key ingredient!)
  - **Scalability**
  - **Self-***: self-organization, self-repair, self-adaptation, self-configuration, self-assembly, etc.
  - **Bottom-up** design and synthesis techniques

Teramac (HP)

- **Approach**: a supervisor locates and tags the defects in the system (crossbar).
- Mapping of an the design on the crossbar.
- **Drawbacks**:
  - slow and impractical if billions of components are involved
  - top-down approach
- **Lessons** (Heath et al., 1998):
  - It is possible to build a powerful computer that contains defects, as long as there is sufficient communication bandwidth.
  - Resources don't have to be regular, but sufficient connectivity is necessary
  - Wires are essential

Christof Teuscher • christof@lanl.gov

NanoFabrics (Seth C. Goldstein & co)

Chemically assembled electronic nanotechnology (CAEN)

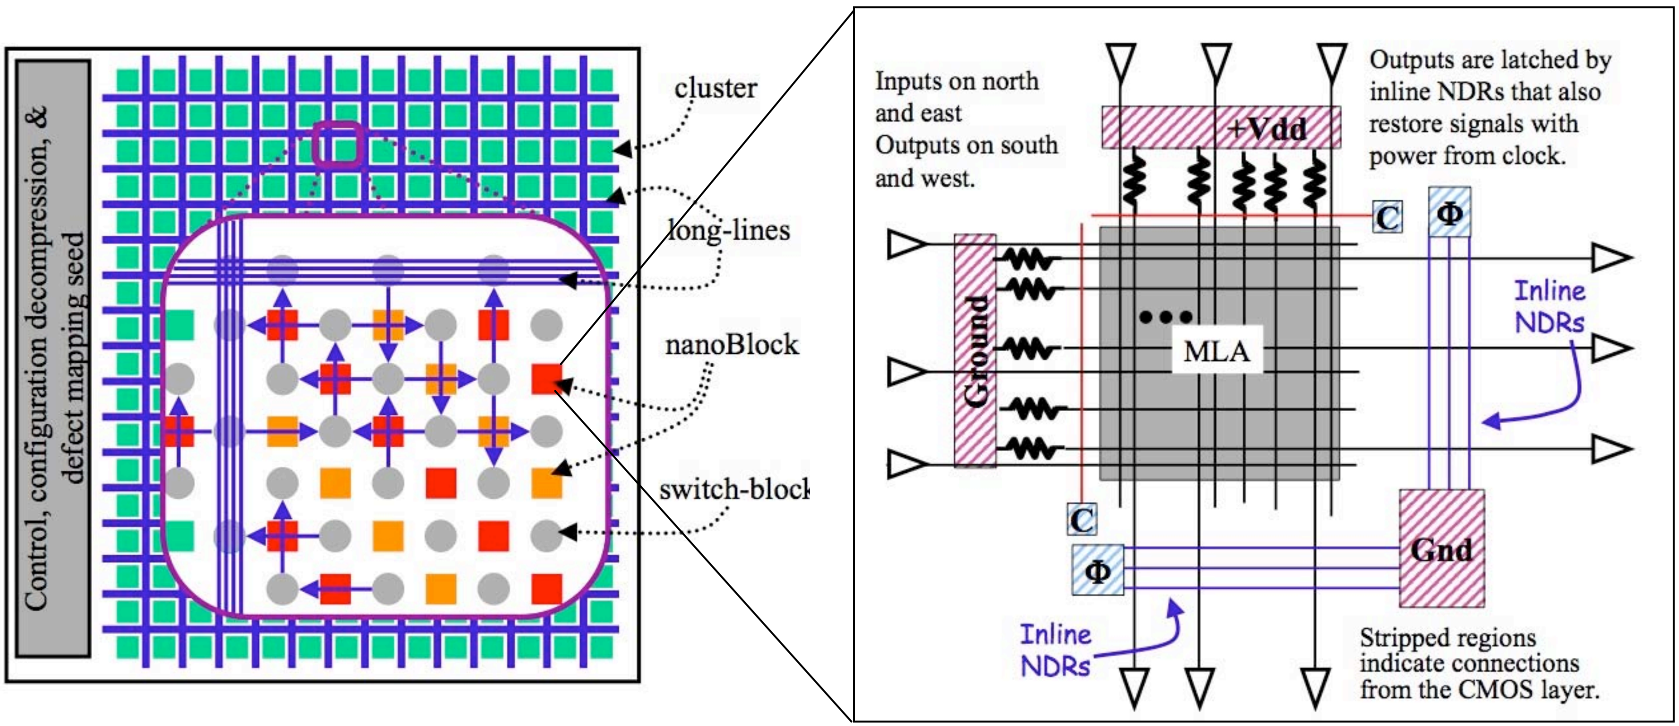

Source: Goldstein, 2001

Christof Teuscher • christof@lanl.gov

Nano Programmable Logic Array (nanoPLA)

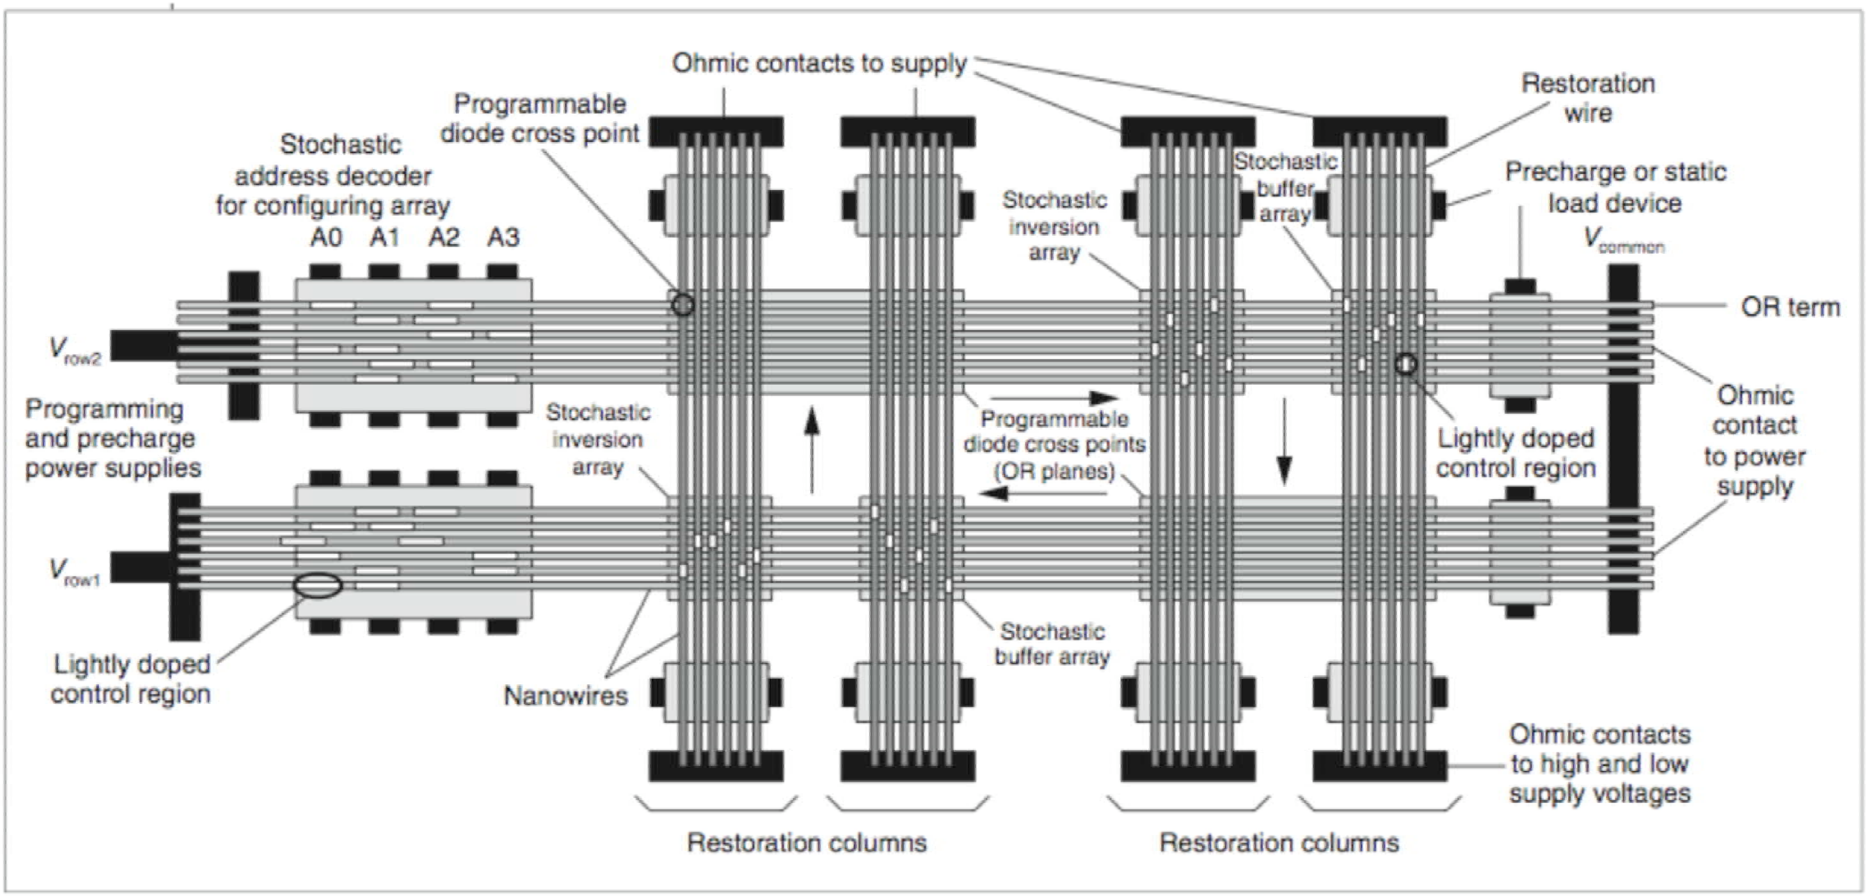

DeHon & Naeimi, IEEE D&T of Computers, 2005

Christof Teuscher • christof@lanl.gov

Nano Programmable Logic Array (nanoPLA) Cluster Tiling with I/O Nanowires

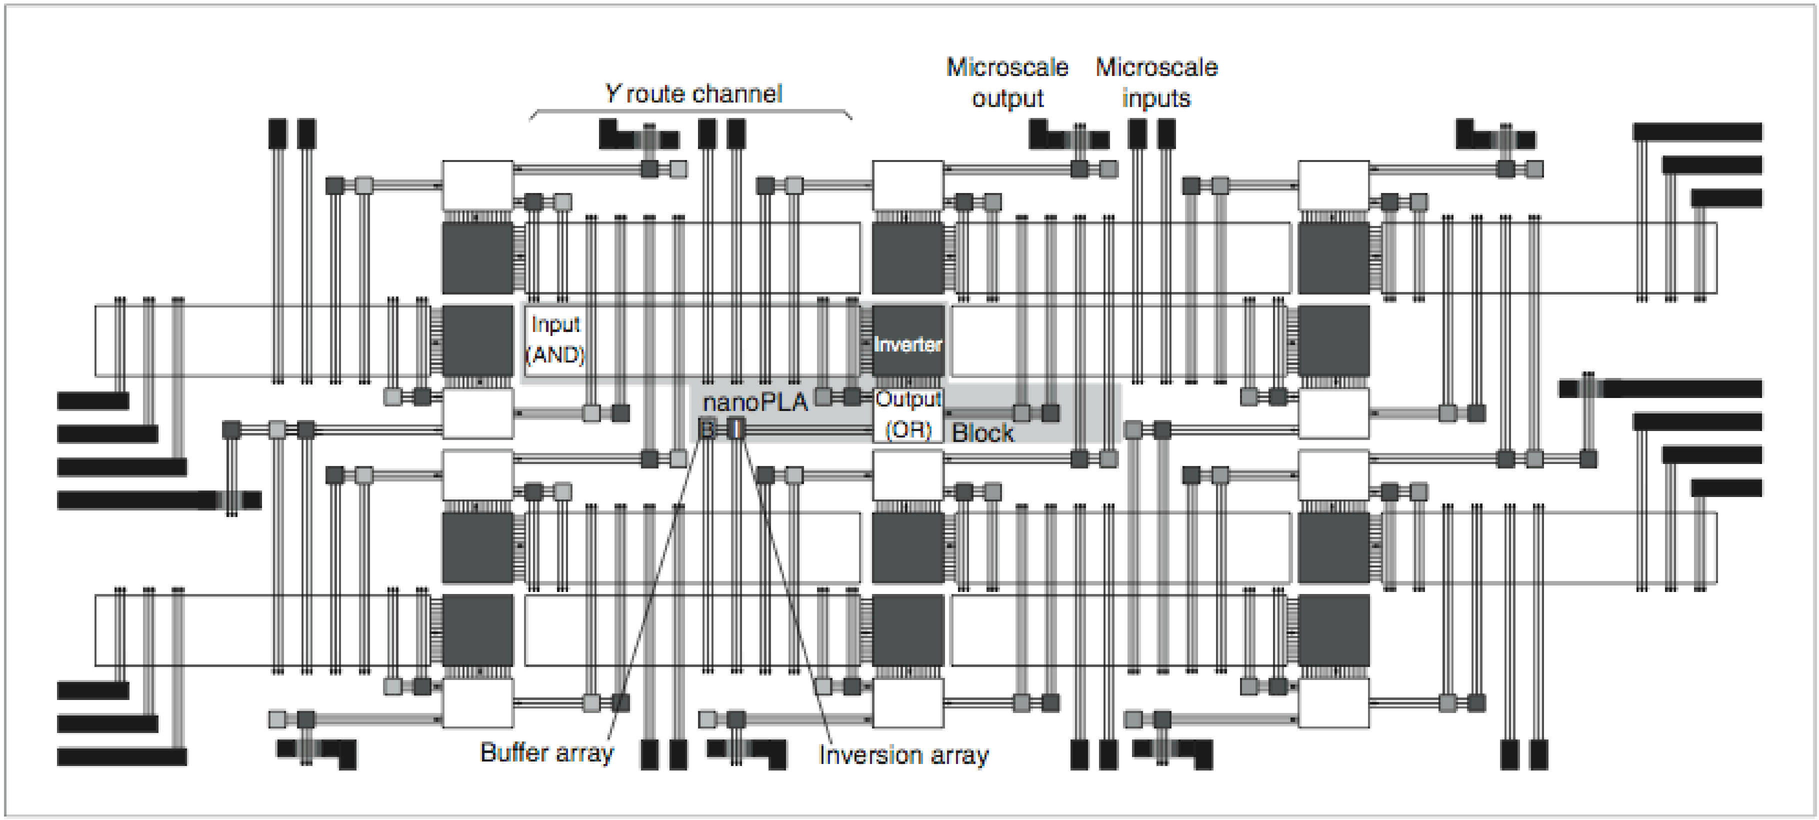

DeHon & Naeimi, IEEE D&T of Computers, 2005

Embryonics Project

- ***Embryonic electronics***: inspired by the embryonic development of living beings.
- Design of highly robust integrated circuits.
- Self-repair and self-healing (cicatrization).
- Self-replication (as an extreme form of self-repair).

- **Toward Robust Integrated Circuits: The Embryonics Approach**. Mange, D.; Sipper, M.; Stauffer, A.; Tempesti, G. In *Proceedings of the IEEE*, vol. 88, num. 4 (2000), p. 516-541
- US patent no 5,508,636, "Electronic System Organised as an Array of Cells", April 16, 1996.

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Christof Teuscher • christof@lanl.gov

Embryonics Landscape

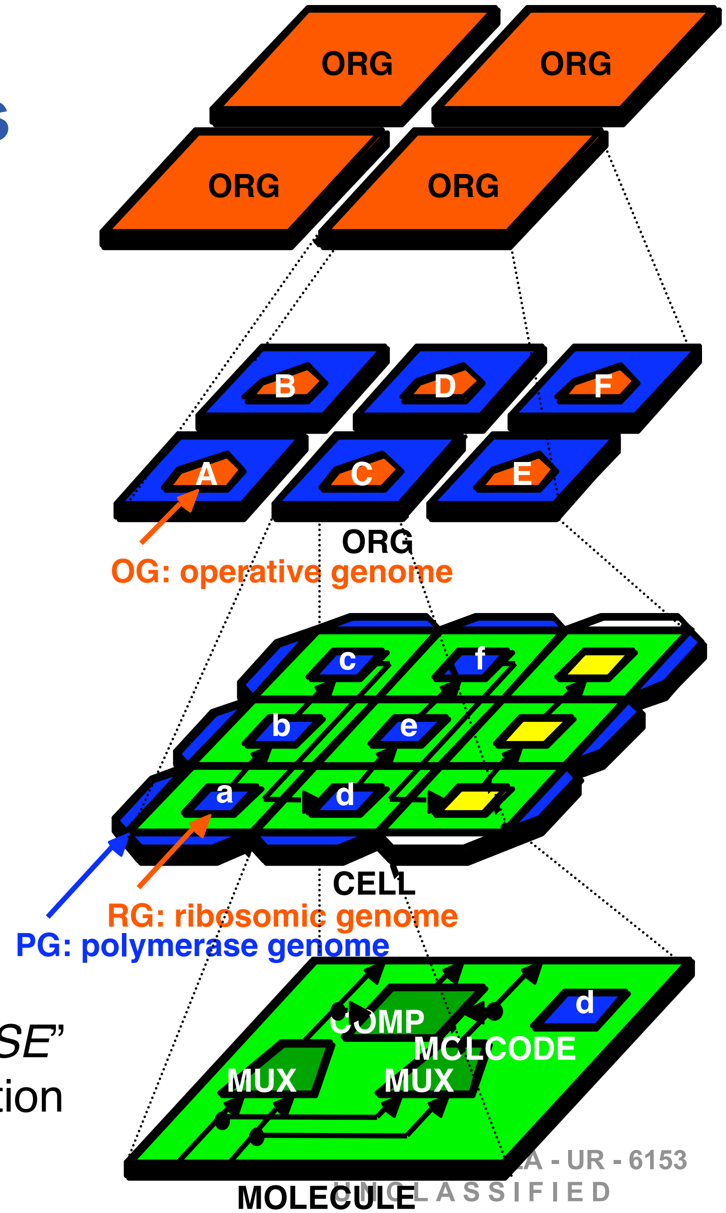

Population level
(population = ∑ organisms)

Organismic level
(organism = ∑ cells)

Cellular level
(cell = ∑ molecules)

Molecular level
(basic FPGA element, MUXTREE)

"IF…THEN…ELSE" instruction

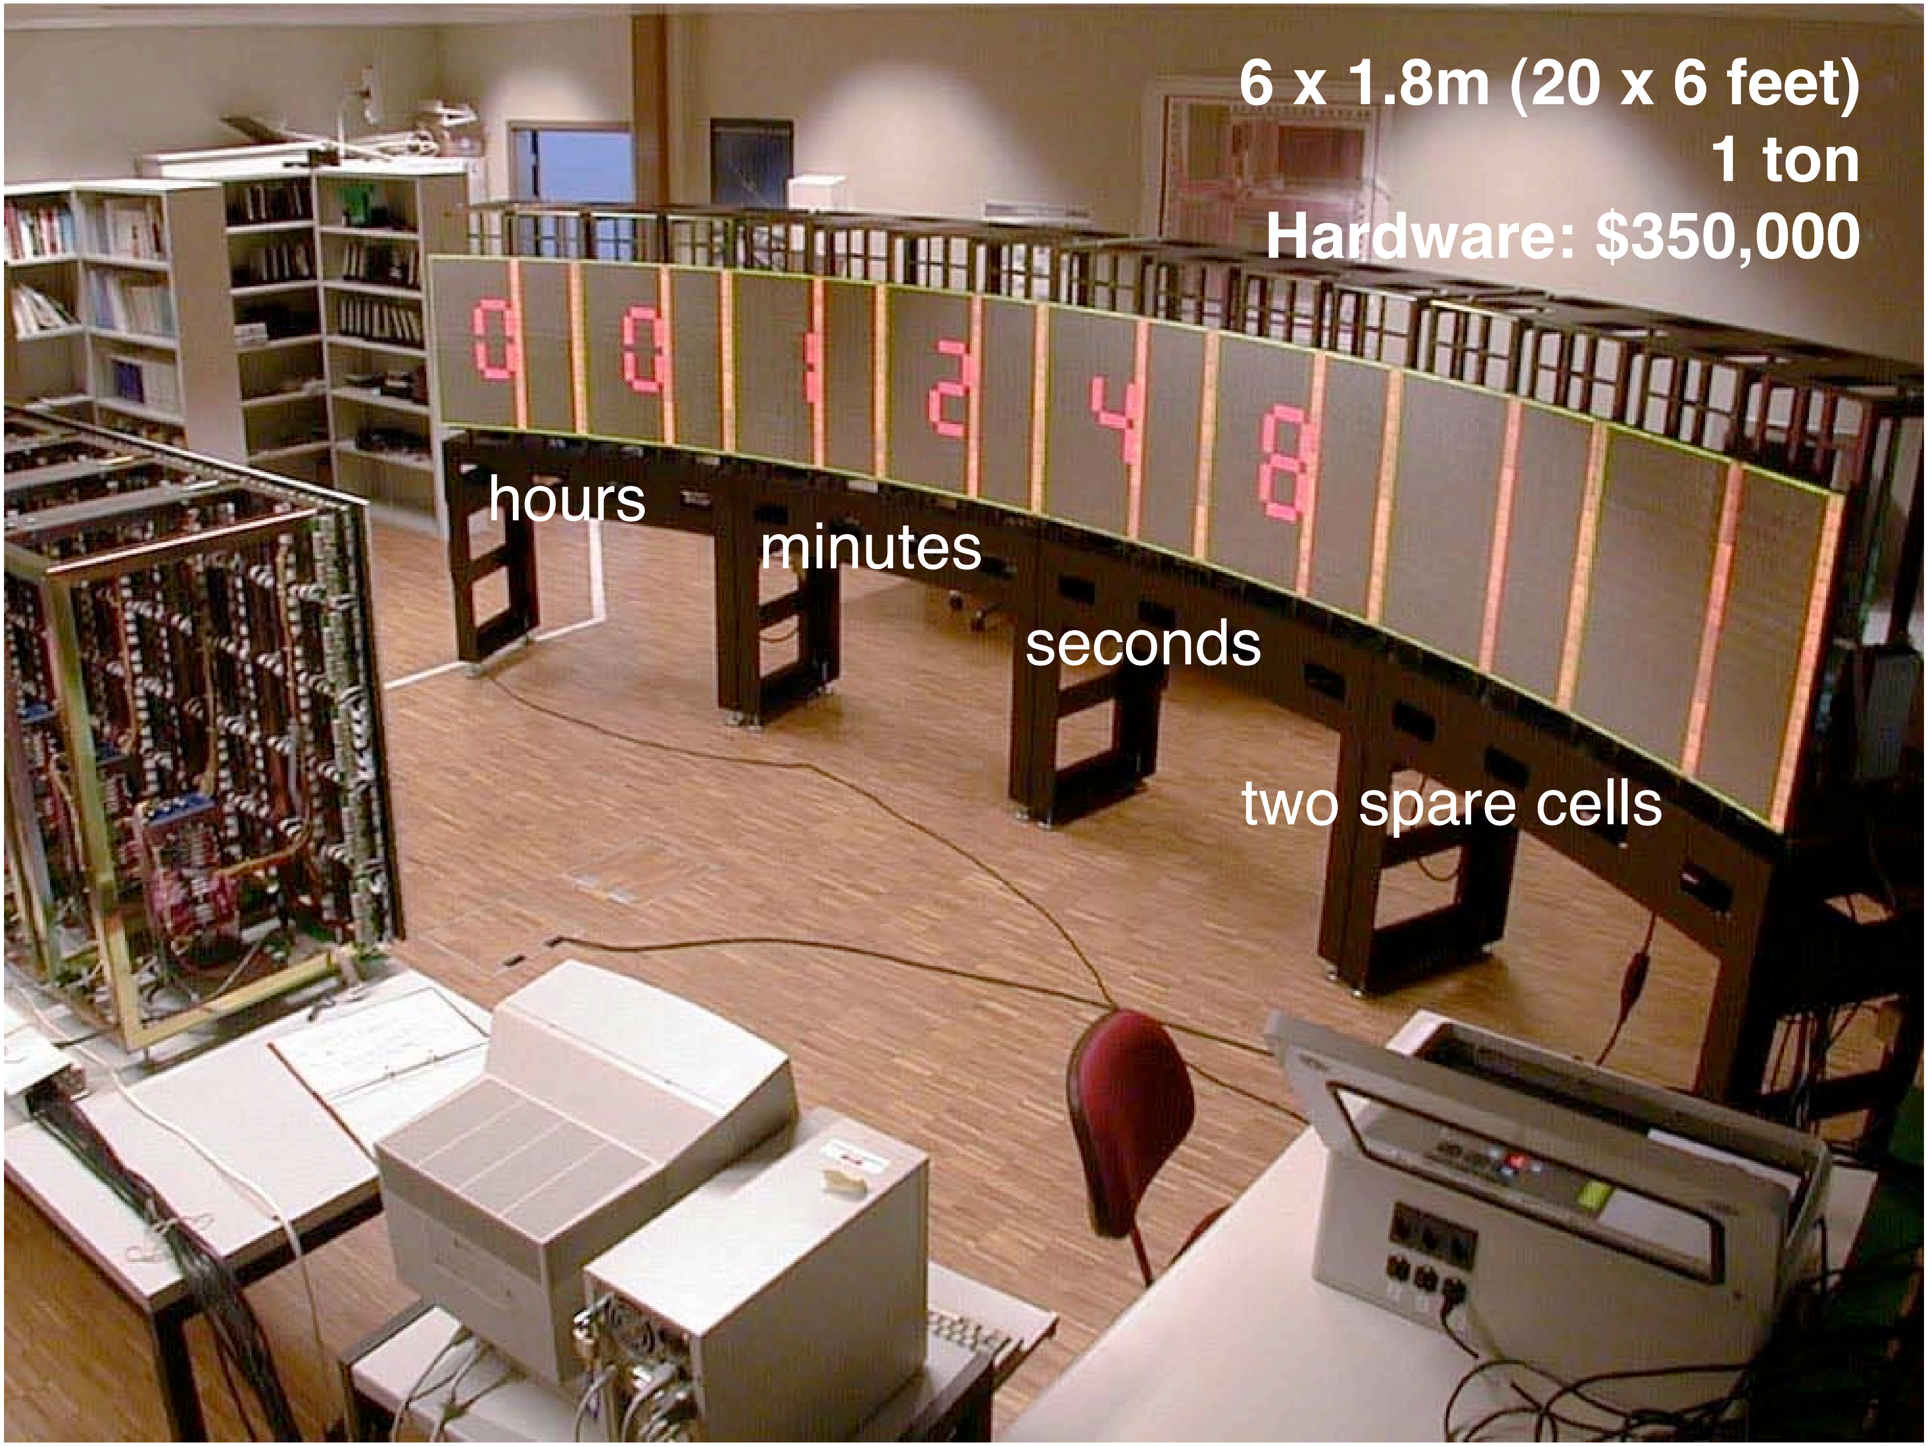
6 x 1.8m (20 x 6 feet)
1 ton
Hardware: $350,000
hours
minutes
seconds
two spare cells

Christof Teuscher • christof@lanl.gov

BioWatch Example: Repair

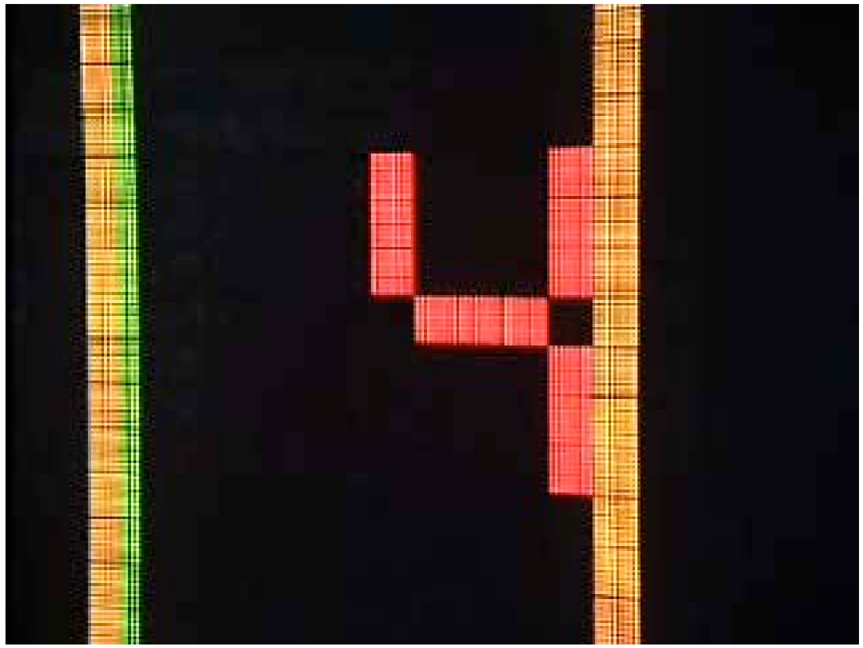

1. Molecular repair
2. Cellular repair (time is saved)

Christof Teuscher • christof@lanl.gov

Tour's Nanocell

- "Create functionality from disorder"
- Molecular states: ON or OFF only

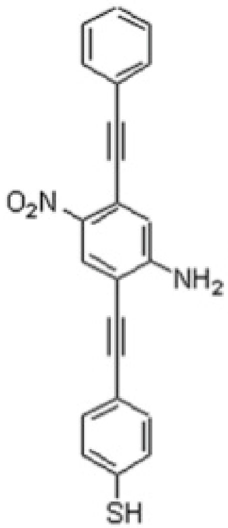

1

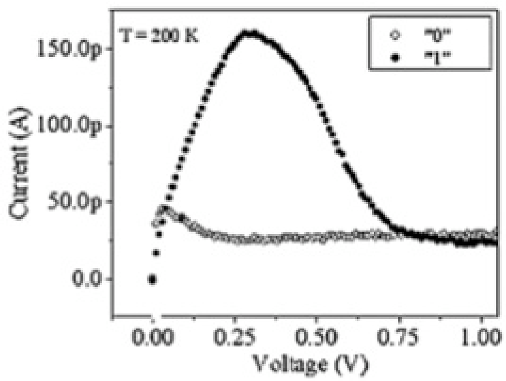

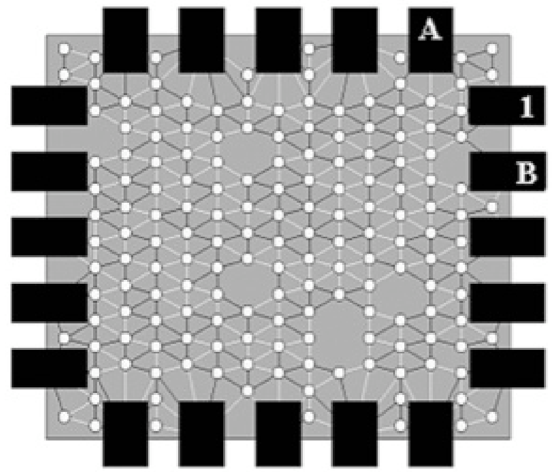

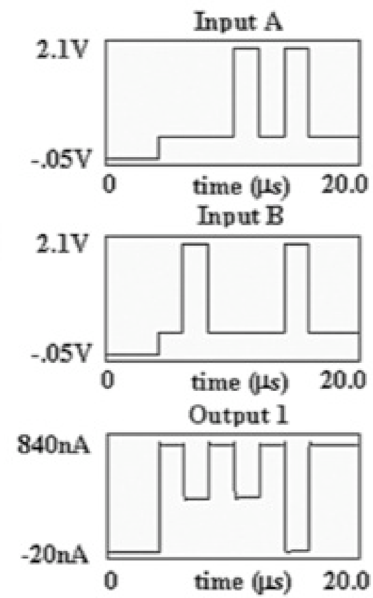

Self-assembled monolayer

Summary of What's Going On

- Fault-tolerance is addressed by:
  - N-tuple (cascaded) modular redundancy
  - NAND multiplexing (von Neumann, 1955)
  - Reconfiguration
- A **regular** (and sometimes even **homogeneous**) array of building blocks is assumed (Exception: Tour's nanocell).
- Hardware complexity is moved to software (e.g., defect avoidance by testing and re-compilation of software)
- No one really knows what the **defect and fault rates** are going to be.

My Research Focus

- **Overall goal:** build faster and more reliable computers with tomorrow's nano materials.

  → we expect a quantum leap in performance

- Four different parts:
  1. Theory
  2. Languages and tools
  3. Simulations
  4. Material

- **Insight:** we need novel (and probably unconventional) paradigms to address the challenges!

My Research Focus

- **Theory**:
  - Reliably compute in **inhomogeneous**, **irregular**, and **uncertain** media.
    - E.g.: Random Boolean networks perform better than cellular automata for global tasks, (Mesot & Teuscher, Physica D, 2005, to be published).
  - Investigate novel **interconnection topologies**
  - Dynamical **hierarchies** to master complexity
  - All aspects of the **self-*** → "lose control in order to gain control"

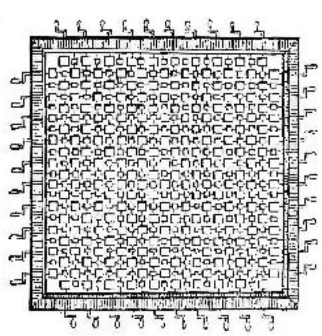

From regularity to irregularity

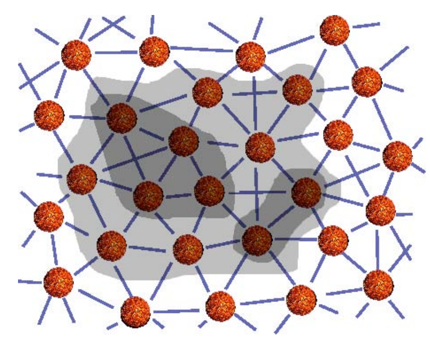

Christof Teuscher • christof@lanl.gov

Trends in Connectivity (Keyes, 2001)

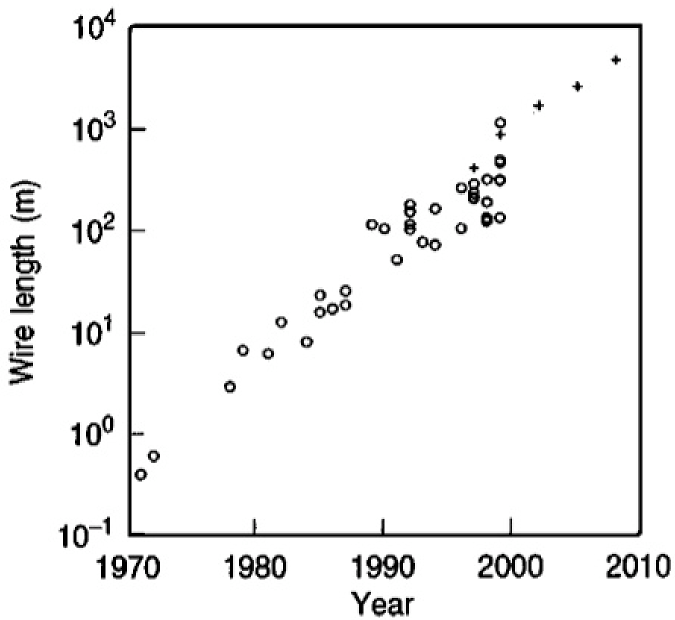

Wire length on μP chips as a function of time. +'s are roadmap projections.

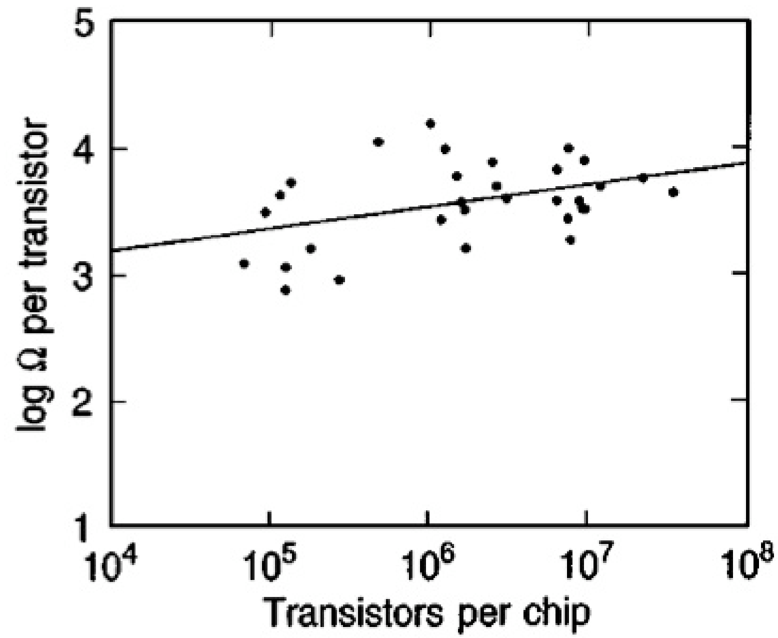

Connectivity Ω per transistor on μP chips according to Meindl, 1995.

Christof Teuscher • christof@lanl.gov

Network on Chip (NoC) Architectures

Pande et al., IEEE Trans. On Comp, 54(8), 2005

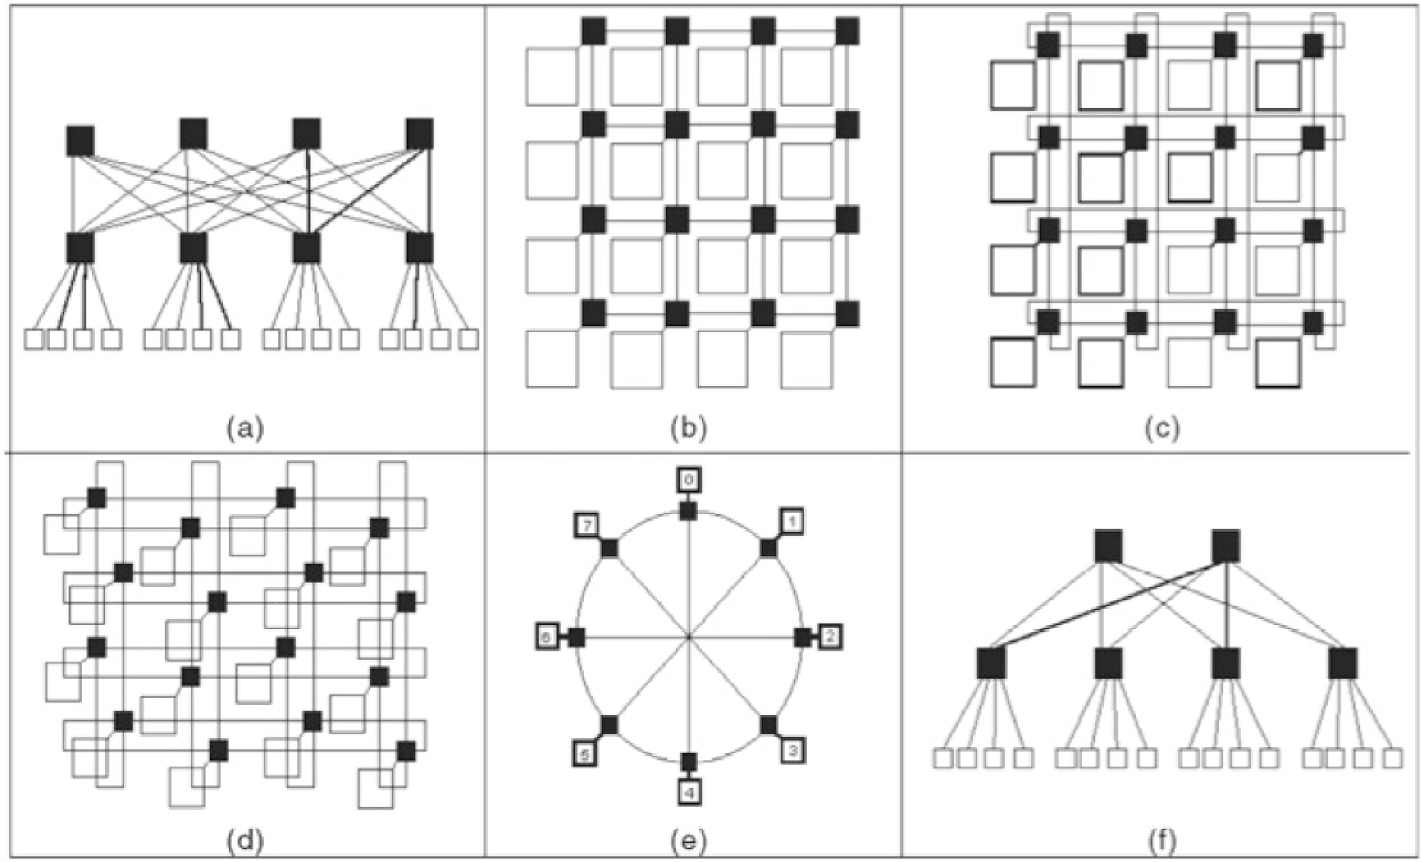

Fig. 1. NoC architectures. (a) SPIN, (b) CLICHÉ, (c) Torus, (d) Folded torus, (e) Octagon, (f) BFT.

Christof Teuscher • christof@lanl.gov

Biologically Inspired Interconnection Architectures: Go Small-World and Scale-Free!

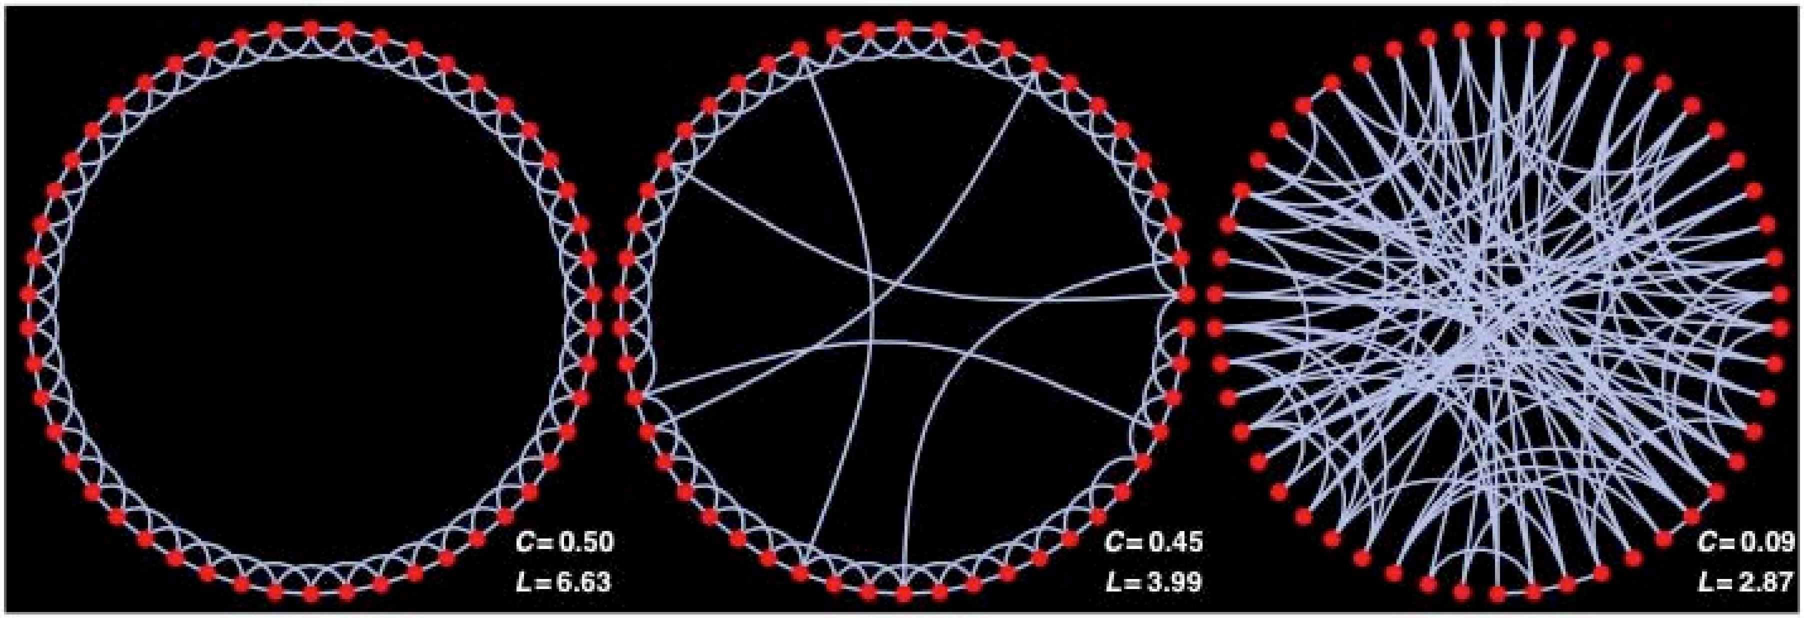

Figure 1. Watts-Strogatz model interpolates between a regular lattice *(left)* and a random graph *(right)*. Randomly rewiring just a few edges *(center)* reduces the average distance between nodes, L, but has little effect on the clustering coefficient, C. The result is a "small-world" graph.

Christof Teuscher • christof@lanl.gov

Amorphous Computing (MIT)

- Myriads of identical elements
- Local broadcast
- Unreliable
- No *a priori* knowledge of their location and orientation.
- Thy cannot be individually programmed.
- Same program for all.

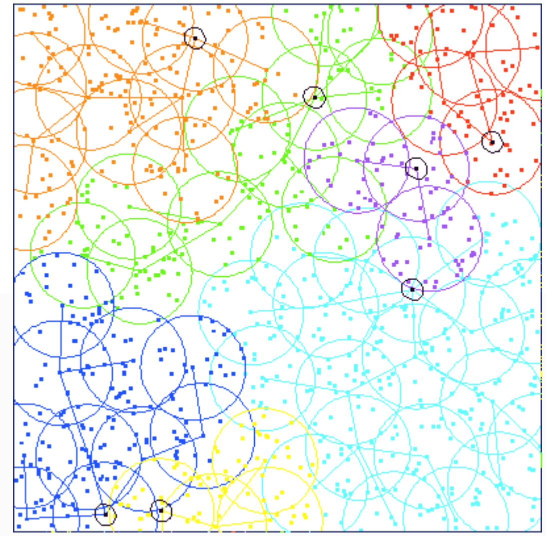

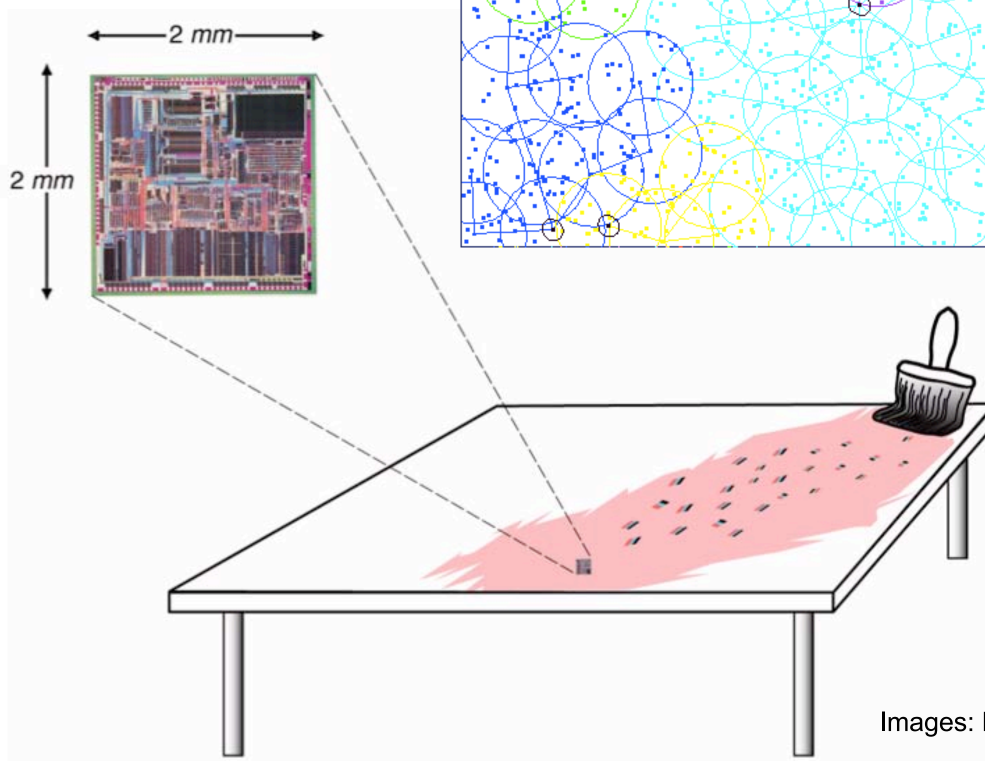

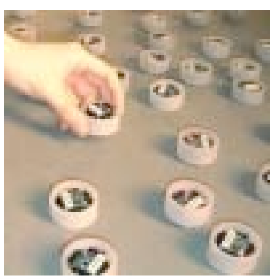

Images: MIT

Christof Teuscher • christof@lanl.gov

Surface Bus on a "Paintable" Computer

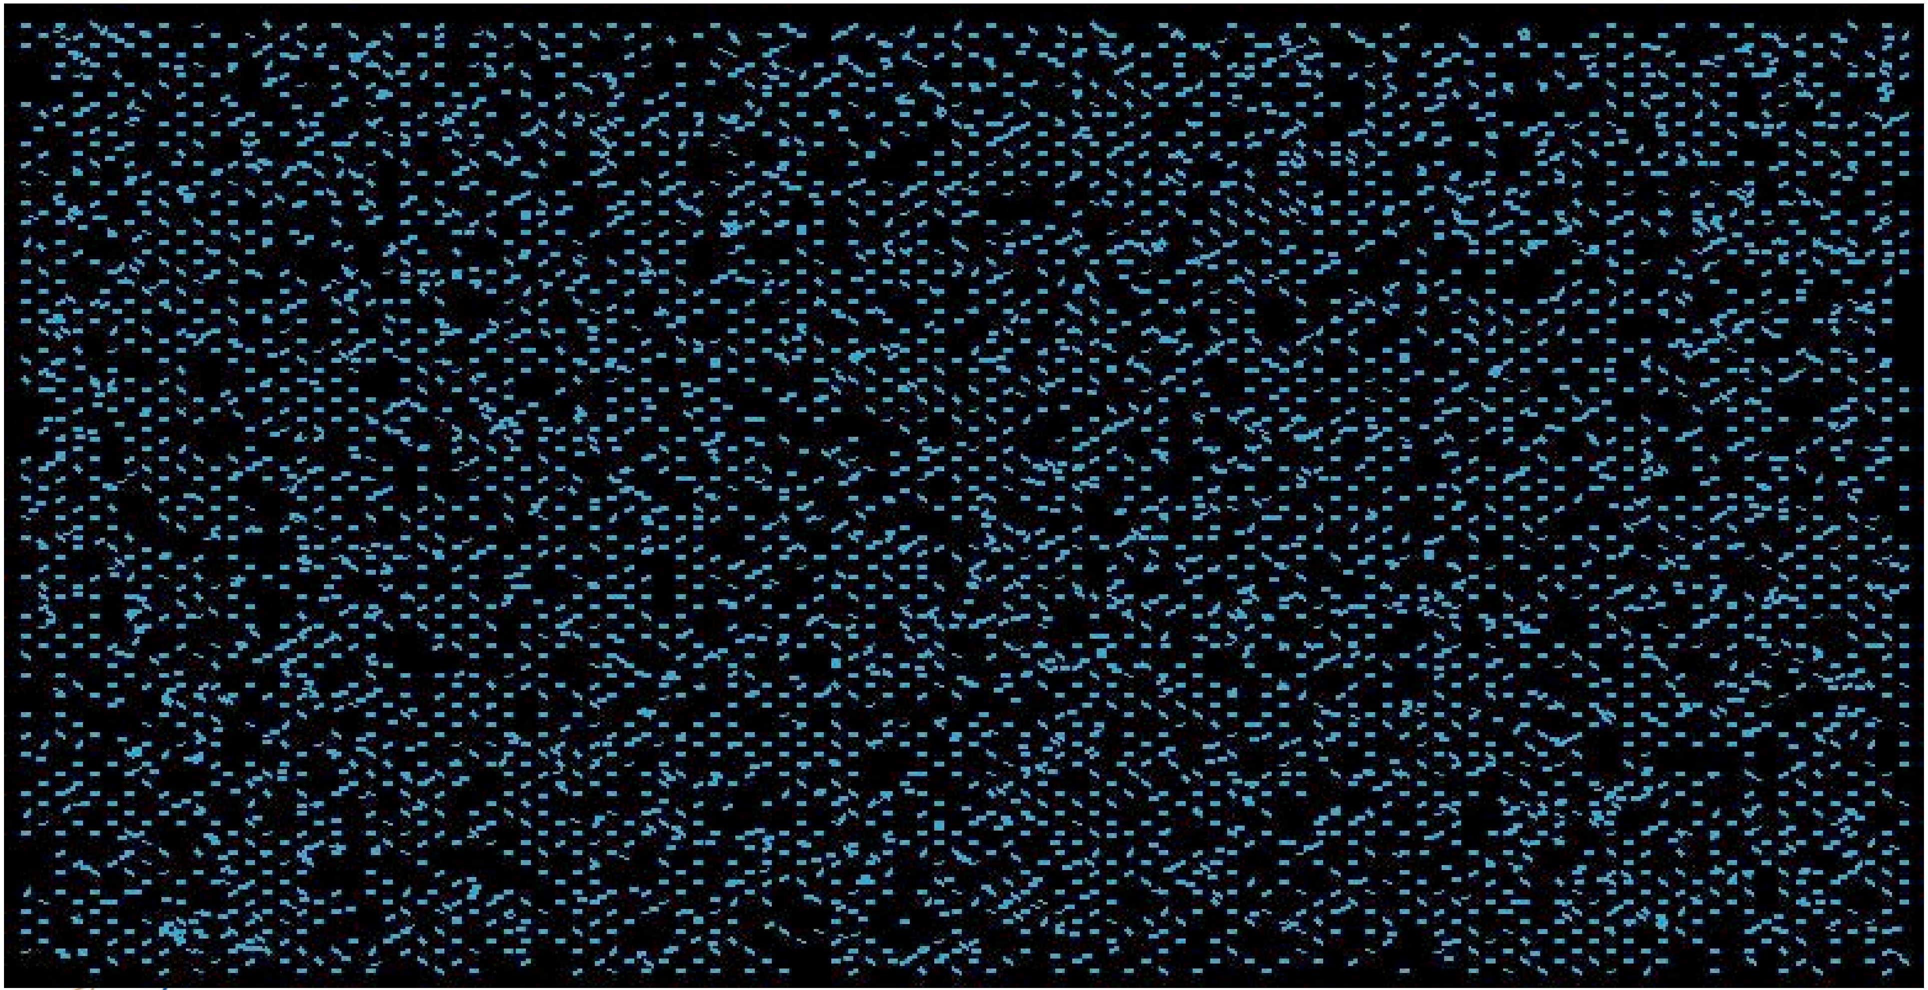

Source: MIT AC website, thesis William J. Butera.

Christof Teuscher • christof@lanl.gov

Dynamical Hierarchies with Membrane Systems (P systems, artificial chemistries)

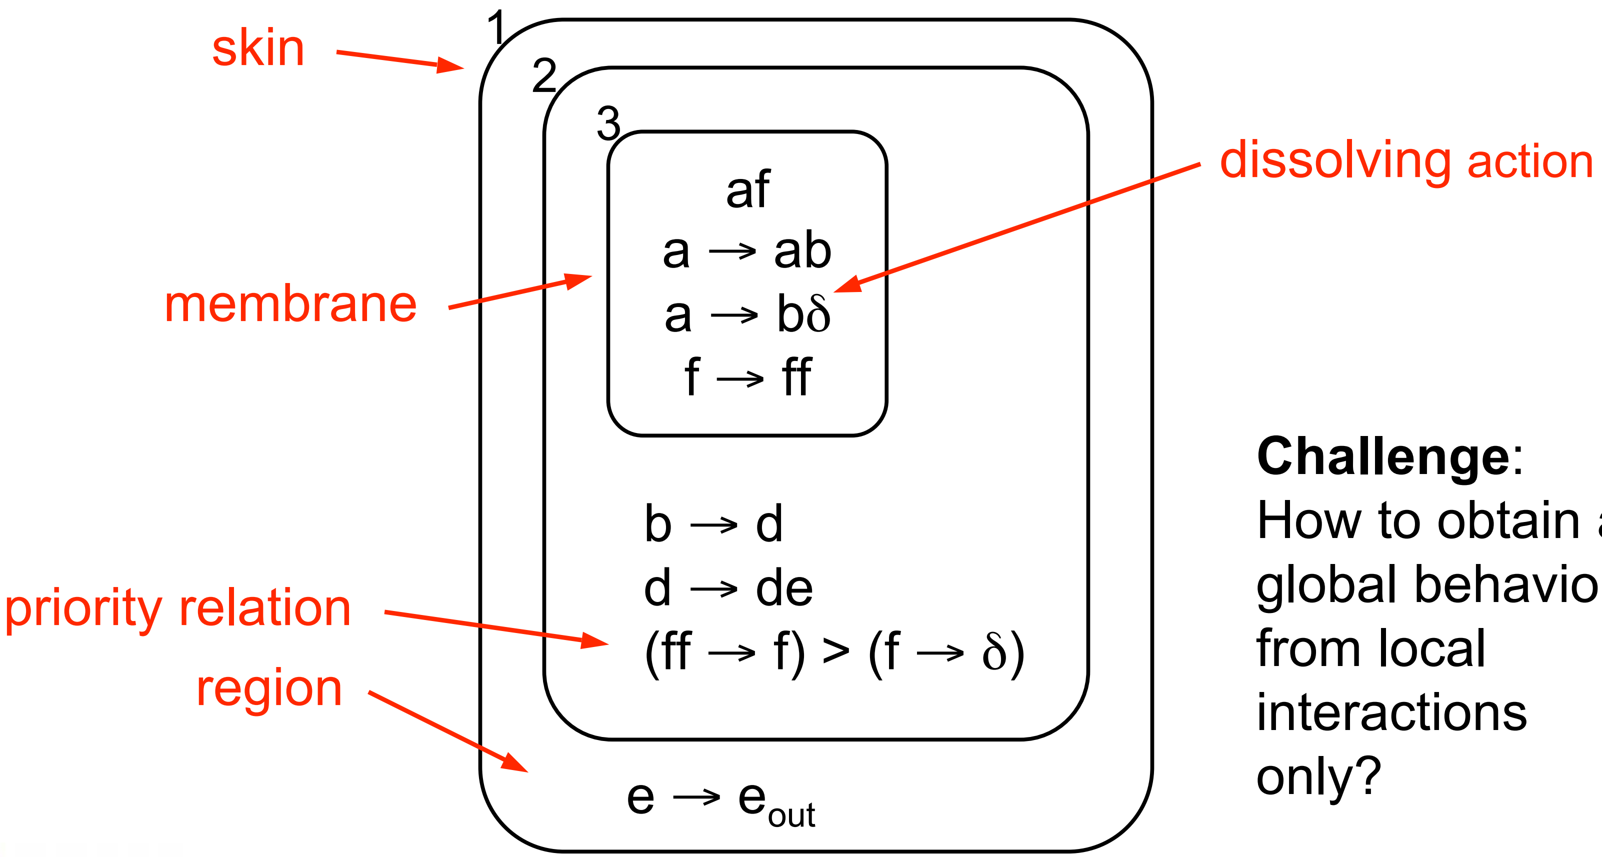

Challenge:
How to obtain a global behavior from local interactions only?

Christof Teuscher • christof@lanl.gov

Programmable Reactor Multitude (PRM)

Example:

- 2D particle arrangement
- Random neighborhood with average connectivity (redundant and dense)
- Particles and links don't move, but they are unreliable

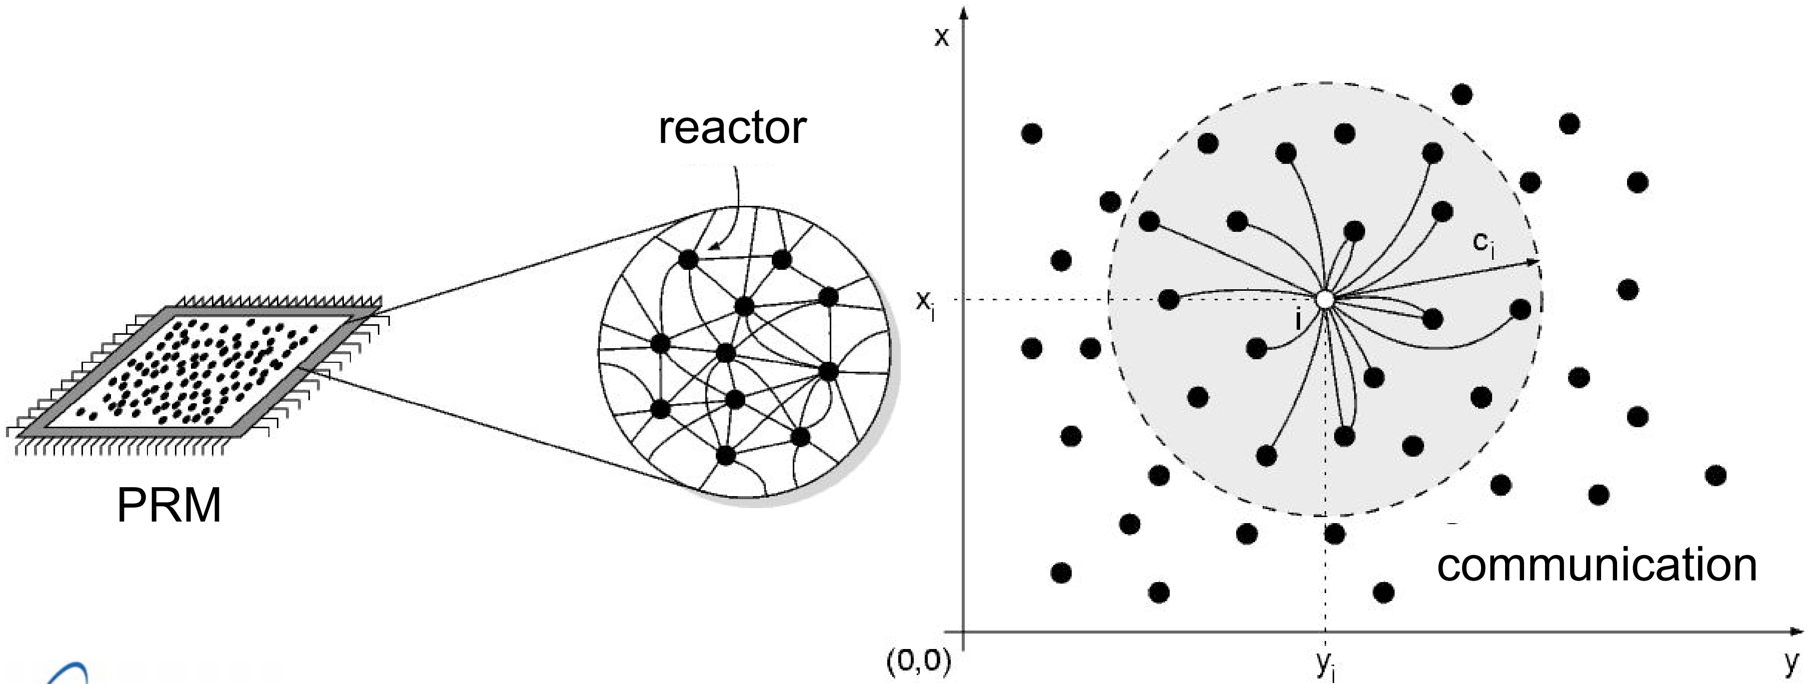

Christof Teuscher • christof@lanl.gov

State Machine in an Artificial Chemistry

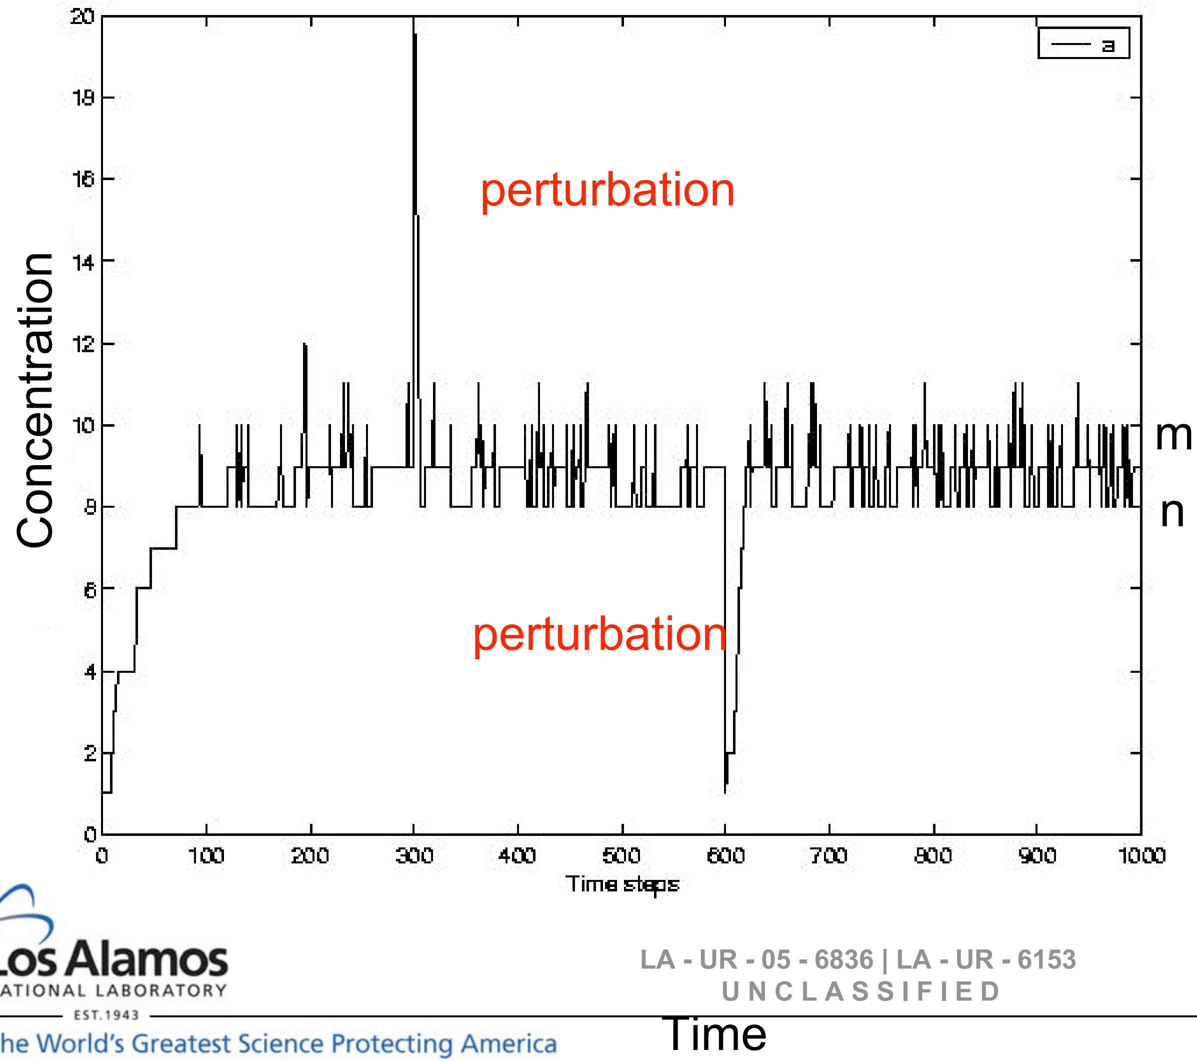

$r = ag \rightarrow c$

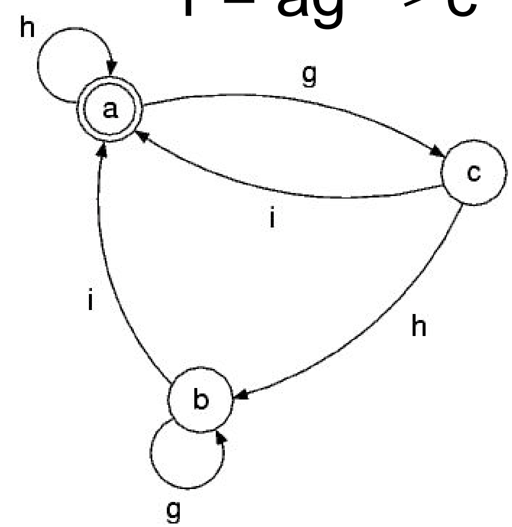

$r_1 = a \rightarrow a^2$

$r_2 = a^m \rightarrow a^n$

My Research Focus (cont.)

- **Languages and Tools:**
  – We need powerful **languages** to program and organize nanoscale systems.
  – Inspiration from *Amorphous computing*, *Blob computing* (F. Gruau), *P systems*, *artificial chemistries*, etc.

  – Conrad's Tradeoff Principle: No system can be at once highly **structurally progammable**, highly **evolutionary efficient** (i.e, adaptable), and highly **computationally efficient**. (Michael Conrad, "The brain-machine disanalogy," BioSystems, 22:179-213, 1989.)

My Research Focus (cont.)

- **Simulations:**
  - **Modeling and simulating** a possible nanoscale architecture from the device to the system.
  - We need to know defect and fault rates, etc.
- **Material:**
  - We're not material scientists!
  - Strong LANL internal and external collaborations
  - Build on top of existing research
  - Ultimately: build a real large-scale system

…So Where's the Innovation?

- Inhomogeneous and irregular substrate → cheaper at fabrication
- No need for testing anymore → cheaper at fabrication
- No global controllers
- Complexity not moved to software
- No global signals

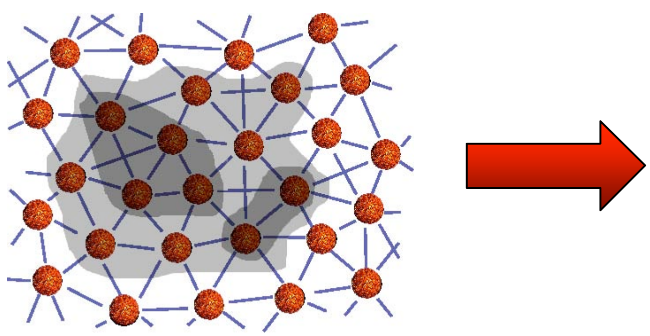

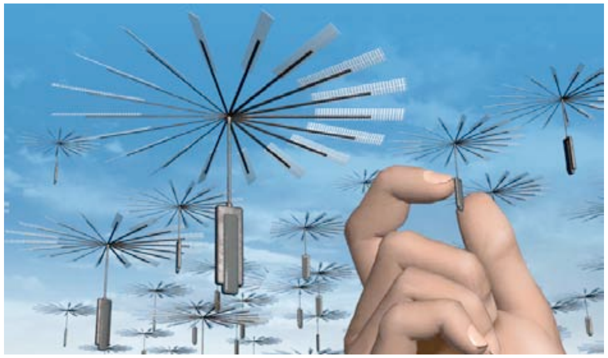

Why LANL? What can we do better here?

- We've got top and diverse researchers on site for this highly **interdisciplinary** undertaking. → Computer science, chemistry, material science, physics, nanotech, …
- We can put things together through **all levels**, from the device to the application.
- We do have the "killer apps."

"The computer revolution hasn't happened yet!"

— Alan Kay
1998